

Facet Classification of Blogs: Know-Center at the TREC 2009 Blog Distillation Task

Elisabeth Lex, Michael Granitzer, and Andreas Juffinger

Know-Center Graz, Inffeldgasse 21a, 8010 Graz, Austria
{elex,mgrani,ajuffinger}@know-center.at

Abstract. In this paper, we outline our experiments carried out at the TREC 2009 Blog Distillation Task. Our system is based on a plain text index extracted from the XML feeds of the TREC Blogs08 dataset. This index was used to retrieve candidate blogs for the given topics. The resulting blogs were classified using a Support Vector Machine that was trained on a manually labelled subset of the TREC Blogs08 dataset. Our experiments included three runs on different features: firstly on nouns, secondly on stylometric properties, and thirdly on punctuation statistics. The facet identification based on our approach was successful, although a significant number of candidate blogs were not retrieved at all.

1 Introduction

On the web, a huge amount of blogs is available. However, the quality of these blogs is questionable in respect to quality because almost no mechanisms exist to validate the content before it is posted online [1]. In this year's TREC blog track, specifically the blog distillation task, exactly this problem is addressed. The goal of the track is to rank candidate blogs that are relevant to a topic or query and to assign the retrieved blogs to a set of facets that represent different quality aspects ¹.

We consider this as a two-step procedure: (i) the first step is to retrieve blogs that are relevant to a topic or query, respectively. This is an Information Retrieval (IR) process. The second step (ii) is to assign particular facets of interest to the retrieved blogs. For the second step, we applied supervised classification on a manually labelled subset of the TREC Blogs08 dataset. Supervised classification has been successfully applied to text classification [2] and facet classification like for instance emotion classification [3]. We trained a classifier on a subset of blogs randomly taken from the TREC Blogs08 dataset² and labelled them manually according to the facets. The classifier was then used to categorise the blogs into the facets.

The rest of this paper is structured as follows: Section 2 outlines the collection and the preprocessing of the data, Section 3 describes the methods and runs, in Section 4, the results are given, and Section 5 concludes our work.

¹ <http://ir.dcs.gla.ac.uk/wiki/TREC-BLOG>

² http://ir.dcs.gla.ac.uk/test_collections/blogs08info.html

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE NOV 2009		2. REPORT TYPE		3. DATES COVERED 00-00-2009 to 00-00-2009	
4. TITLE AND SUBTITLE Facet Classification of Blogs: Know-Center at the TREC 2009 Blog Distillation Task				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Know-Center Graz,Inffeldgasse 21a,8010 Graz, Austria,				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Proceedings of the Eighteenth Text REtrieval Conference (TREC 2009) held in Gaithersburg, Maryland, November 17-20, 2009. The conference was co-sponsored by the National Institute of Standards and Technology (NIST) the Defense Advanced Research Projects Agency (DARPA) and the Advanced Research and Development Activity (ARDA).					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 6	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

2 Collection and Preprocessing

The experiments for the blog distillation task are carried out on the Blogs08 dataset. The Blogs08 dataset has 28,488,767 blog posts from 1,303,520 blog feeds³. It samples the blogosphere from January 2008 to February 2009.

For the first step, the retrieval task, we vectorised and indexed 606,939 different blogs out of 1,303,520 blogs. From our indexed blogs, we retrieved 3,476 million different blog feeds. Our search index size was 41 GB, and the size of the whole repository with vectors for all features was 220 GB. Previous experiences have shown that an analysis has to be performed on blog post level [6]. To extract the blog posts, we exploited the permalinks.

For the second step, the classification task, we manually annotated 83 blogs by the facets given in the TREC blog task. These 83 blogs were selected randomly to guarantee an independent sample. Also, due to the random selection, we assume that the label distribution reflects the real blog world although the number of blogs is not high. As mentioned before, we assigned the facets on blog level. However, our analysis was done on blog post level, so if the whole blog was assigned e.g. the facet “opinionated”, we assigned this label to all its blog posts. This resulted in 12870 annotated blog posts. Since we did not check all the labelling of all 12870 blog posts, there might be some mislabelled items. We then transformed the unstructured blog posts in a structured form and thus prepared the blog posts for our text mining unit. Note that we first filtered the blogs by language in order to retrieve only English blogs. For this, we exploited a language guesser developed at our institution that is based on n-grams and the Apache Nutch project⁴. From the remaining structured blog posts, we extracted nouns, sentences, and punctuation features using an information extraction tool developed at our institution based on OpenNLP⁵. For each set of features, we created a vector and stored these vectors in our index.

Our search framework was based on Apache Lucene which is an open source search engine that is able to index several gigabytes of documents at reasonable time. Also, complex searches can be performed quickly on the documents. Besides, Lucene maintains a relevance ranking scoring which is applicable to find the best matches to queries. We used this relevance score for ranking, as described later. To sum up, our resulting training set consists of 12870 blog posts. We trained our classifier on the annotated blog posts and also classified the blogs on blog post level. To assign the facets to the whole blogs, we performed a majority voting over the posts.

³ http://ir.dcs.gla.ac.uk/test_collections/blogs08info.html

⁴ <http://lucene.apache.org/nutch/>

⁵ <http://opennlp.sourceforge.net/>

2.1 Relevance Ranking

In order to derive a final relevance ranking, we normalised the Lucene score so that it lies between 0 and 1 by dividing it by the maximum score which is shown in Equation 1.

$$luceneScore = score / maxScore \quad (1)$$

Then, we calculated the final relevance score whereas we took into consideration both the Lucene score and the classifier confidence for the particular facet. The formula for the final score is shown in Equation 2.

$$finalScore = \alpha * luceneScore + (1 - \alpha) * facetConfidence \quad (2)$$

3 Methods and Runs

For the retrieval task, we used Apache Lucene as search engine, as mentioned before. With Lucene, we searched our blog index for the 50 topics and ranked and sorted the top 100 blogs according to our relevance ranking scheme (see Equation 2). We then classified the top 100 blogs into the given facets using the trained classifier. As a classification algorithm, we applied a Support Vector Machine based on LibLinear [4] with standard parameterisation.

We carried out three runs whereas each run is characterised by the features used in this run:

1. *Nounfull NF*: For this run, we used a feature space only with nouns annotated by the OpenNLP library⁶.
2. *Punctfull PF*: In this run, we created a feature space only with punctuation features. Punctuation features are for instance the amount of double or triple punctuations.
3. *Sentencefull SF*: For this run, we used a feature space with stylometric properties based on sentence statistics. Examples for such features are the average number of words per sentence, the average number of unique POS tags per sentence, the ratio of lower case characters to upper case characters, the number of adjectives or adverbs per tokens.

Clearly, a number of different features can be used for text classification problems, as described in [7]. In our feature selection strategy, we considered both lexical and stylometric features. Nouns belong to the category of lexical features and generally cover topics. Usually, a classification based on nouns yields good results, at least in our experience. Punctuation features and features based on sentence statistics are stylometric features. In general, stylistic variations are described by stylometric features and therefore such features are often used for authorship identification [5]. In our experiments we applied stylometric features to guarantee topic-independence. Our goal was to avoid that the classifier learns topics instead of characteristics.

⁶ <http://opennlp.sourceforge.net>

4 Results

With our methodology, we achieved the following results:

Of the 50 topics, we were able to retrieve 39 topic which had at least one relevant blog for each side of the facet (e.g., one relevant opinionated blog and one relevant factual blog). Therefore, all our results relate to these 39 topics. Note that we did not manage to vectorize and index the whole Blogs08 dataset⁷. Therefore, relevant blogs were missing in our index which is reflected in poor results for certain topics. In Table 1, the results for all runs are summarised. For each run, three rankings of the blogs were performed. The first, “none”, denotes a ranking with no facet value applied and serves therefore as baseline. The second “first” correspond to the first value of a facet (“opiniated”, “in-depth”, “personal”) and the third “second” represent the second value of a facet (“factual”, “shallow”, “official”).

RUN	Facet	MAP	R-Precision	P@10
NF	none	0.0624	0.0980	0.1410
NF	first	0.0538	0.0725	0.0769
NF	second	0.0231	0.0346	0.0256
PF	none	0.0624	0.0980	0.1410
PF	first	0.0691	0.0846	0.0692
PF	second	0.0227	0.0339	0.0308
SF	none	0.0624	0.0980	0.1410
SF	first	0.0559	0.0717	0.0667
SF	second	0.0302	0.0334	0.0436

Table 1. Results for all facets for all runs

From Table 1 one can see that the obtained MAP values for all rankings are quite low. In general, MAP corresponds to the average precision and emphasises returning more relevant documents earlier. Due to the fact that we were only able to index less than half of the data, our retrieval was suboptimal. More specifically, in many cases, we did not retrieve enough relevant blogs for a topic and this lead to quite low MAP values. For instance, in run NF with the second facet enabled, only 69 blogs out of 331 relevant blogs were retrieved and consequently, the MAP value for this run was low, as shown in Table 2.

In cases where our retrieval was successful which means we obtained enough relevant blogs for a topic, also good MAP results were achieved. An example of this is outlined in Table 3, which holds the results for the first facet for topic 1101.

Another example where the classification was successful was topic 1103, as shown in Table 4.

⁷ Unfortunately, the lucene index broke a few days before the deadline and we had to re-index the data.

	Results
numQueries	39
numRetrieved	3801
numRelevant	331
numRelevantRetrieved	69
map	0.0231

Table 2. NF: Results for second facet - all

TOPIC	RUN	MAP	R-Precision	P@10
1101	NF	0.2590	0.4000	0.5000
1101	PF	0.2989	0.4571	0.6000
1101	SF	0.2221	0.4143	0.3000

Table 3. SENTENCE: Results for first facet for topic 1101

TOPIC	RUN	MAP	R-Precision	P@10
1103	NF	0.0525	0.1429	0.1000
1103	PF	0.2857	0.2857	0.2000
1103	SF	0.2381	0.2857	0.2000

Table 4. SENTENCE: Results for first facet for topic 1103

Compared to the summary statistics for the first facet ”opinionated” , which is given in Table 5, one can see, that for this topic 1103, our results were above average for the runs ”punctfull” and ”sentencefull”.

MAP1.	MAP2	MAP3	R-Prec1.	R-Prec2
0.4286	0.1667	0.0000	0.4286	0.2857

Table 5. SENTENCE: Results for first facet for topic 1103

In some cases, our relevance ranking was not ideal as is reflected in low values for R-precision. In later experiments on the Blogs08 dataset where we also used the annotated blogs for training, we were able to achieve a classification accuracy of 0.75 for nouns only. In other experiments on other features e.g. stems, we reached an accuracy of 0.91 on blog post level. A deeper analysis of the results of the facet classification in this work revealed that for the first facets personal, opinionated and indepth, we achieved a classifier accuracy of 80%. Summing up, if we improve our retrieval and ranking, we can definitely improve our results. Also, later experiments revealed that if we use features like stems, an improvement of 15-30% of the classification accuracy can be achieved.

5 Conclusions

In this work, we described the efforts and results for our participation in the TREC blog track 2009. Our system is based on a plain text index extracted

from the XML feeds only. We successfully indexed 680k of 1.3 Mio blogs and this index was used to retrieve candidate blogs for the given topics. From the top 2500 result blog entries, the top 100 blogs were identified according to the accumulated relevance score of the particular blog entries. The resulting blogs were classified using a Support Vector Machine trained on a manually labelled subset of the TREC Blogs08 dataset. Three runs were conducted, one based on nouns, one based on stylometric properties, and one based on punctuation statistics. The facet identification based on this approach was successful, although a significant number of candidate blogs were not retrieved at all. A deeper analysis of our facet classification results revealed that for the first facets personal, opinionated and indepth, we achieved a classifier accuracy of 80%. Our best run was based on sentence statistics for the facet “Personal”: we achieved the fifth best results of all groups with this run. The results for the facet “Opinionated” were our second best results, also based on sentence statistics. Note that sentence statistics are a very simple feature. Clearly this is a very encouraging finding, since apparently very simple stylometric features allow to address certain quality aspects of blogs.

6 Acknowledgements

The Know-Center is funded within the Austrian COMET Program - Competence Centers for Excellent Technologies - under the auspices of the Austrian Ministry of Transport, Innovation and Technology, the Austrian Ministry of Economics and Labor and by the State of Styria.

References

1. C. C.-D. D. G. A. Agichtein, E. and G. Mishne. Findin high-quality content in social media with application to community-based question answering. In *Proceedings of WSDM*, 2008.
2. A. An, L. Viii, L. Viii, L. Viii, T. Joachims, T. Joachims, F. Informatik, F. Informatik, F. Informatik, F. Informatik, and F. Informatik. Text categorization with support vector machines: Learning with many relevant features, 1998.
3. H.-H. C. Changhua Yang, Kevin Hsin-Yih Lin. Emotion classification using web blog corpora. In *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence*, 2007.
4. R. Fan, K. Chang, C. Hsieh, X. Wang, and C. Lin. Liblinear: A library for large linear classification. *Journal of Machine Learning Research*, 2008.
5. J. Grieve. Quantitative authorship attribution: An evaluation of techniques. *Literary and Linguistic Computing*, 2007.
6. A. Juffinger, M. Granitzer, and E. Lex. Blog credibility ranking by exploiting verified content. In *WICOW '09: Proceedings of the 3rd workshop on Information credibility on the web*, pages 51–58. ACM, 2009.
7. J. Karlgren and J. Cutting. Recognizing text genres with simple metrics using discriminant analysis. In *Proceedings of COLING*, pages 1071–1075, 1994.